Transmission – Linux Audio goes mobile

Daniel James
64 Studio Ltd.
P.O. Box 37, Freshwater,
Isle of Wight, PO40 9ZR, Great Britain
daniel@64studio.com

Free Ekanayaka
64 Studio Ltd.
P.O. Box 37, Freshwater,
Isle of Wight, PO40 9ZR, Great Britain
free@64studio.com

Abstract

Transmission is a custom GNU/Linux distribution designed for real-time audio production on mobile computing platforms, including Intel 'Ultra Mobile PC' devices. It was developed on behalf of Trinity Audio Group by 64 Studio Ltd. and was released to the public in the summer of 2008.

Keywords

distribution, mobile, real-time

1 Introduction

Ronald Stewart is a record producer, originally from Seattle, USA but now based in Los Angeles. He was one of the first users of the Digidesign ProTools system, the de facto standard in proprietary DAWs since the early 90's. Back then, Stewart was producing remixes for the likes of cash-burning Discordianist pop duo The KLF. Fast forward to the middle of this decade, and Stewart found himself thinking about a better kind of DAW for his own needs. It had to be very portable, easy to use and with a long battery life. Most importantly, it needed to have first-class audio quality without requiring extra boxes and cables to be carried around.

In so-called Agile software development, projects begin with a user story. Stewart's user story was that of a record producer who does a lot of travelling by air. The computer he had in mind had to be small enough to be used comfortably while sitting in an airline seat, and would be able to do CPU-intensive audio processing for the whole flight, without depending on external power. It would enable him to put down his thoughts in audio form or complete an edit while still in the air, and upload the result to the studio over any Wi-Fi link available after arriving at the airport.

2 Prototype hardware

Stewart's conceptual starting point was Sony's PlayStation Portable games console. He reasoned that a device like this, much smaller than a laptop but with a fast CPU, could feature pro audio software and have an audio interface with quality to match. Since the PSP is not officially open to independent development, Stewart assembled a prototype using a small VIA x86 system board and a six-and-a-half-inch TFT display.

Stewart had heard that GNU/Linux was efficient enough to use on low-power embedded devices, and had multimedia applications available. A number of manufacturers in the pro audio sector had recently launched embedded products based around the Linux kernel, taking advantage of reduced R&D costs and shorter time-to-market, due to the number of high-quality Free Software building blocks already available. These manufacturers gain OS vendor independence into the bargain.

Xubuntu, the variant of Ubuntu with an XFCE desktop, was loaded on the miniature iPod-style hard disc used in Stewart's first prototype. Audacity was used to make some test recordings through an Echo Audio interface, which worked well enough for Stewart to realise he was on to something. He gave his project the code name Trinity, and founded a start-up called Trinity Audio Group to develop the device commercially.

A Compulab CM-iGLX module and development board were ordered for a second prototype. The CM-iGLX is about the size of a credit card, and includes a low-power AMD Geode LX800 CPU clocked at 500 MHz, 256MB of DDR RAM, 512MB Flash storage and many of the interfaces you would expect to find on a full-sized

motherboard. Geode LX CPU's are used by the One Laptop Per Child project; while they are x86 compatible and will run a standard distro, optimisations for these chips have been contributed to the Linux kernel.

Figure 1: Second prototype, 2006.

The CM-iGLX has typical consumer audio features available, including an AC97 codec, but the sound quality requirements of the Trinity product meant that a custom baseboard would have to be fabricated. Also, most studios use condenser microphones which require pre-amplifiers with phantom power. Consumer audio interfaces don't provide phantom power, so this was another consideration which led Stewart to specify a custom baseboard.

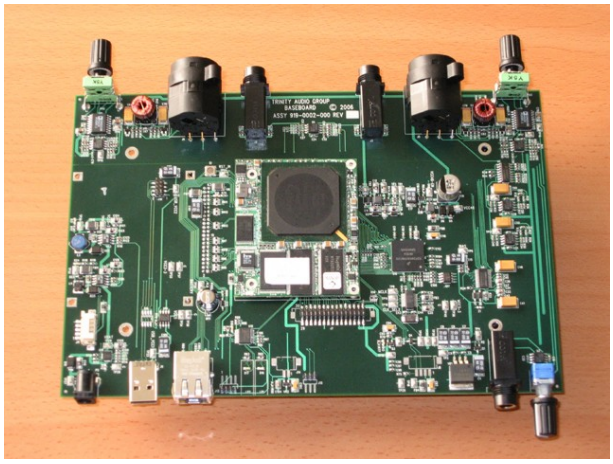

Figure 2: Custom Trinity baseboard with CM-iGLX module in the centre, 2006

A high-quality sound chipset was licenced from Echo Audio and incorporated into the baseboard design. The CM-iGLX module fitted into a socket on the baseboard, which also provided Neutrik combination microphone/line inputs with level trim potentiometers, line output jacks, a headphone output with level control, and standard external USB and power connectors.

3 Into production

The first Trinity pre-production model, consisting of the custom baseboard and other parts packed between two sheets of Plexiglas, was demonstrated in public for the first time at the 5th International Linux Audio Conference in Berlin, Germany during March 2007.

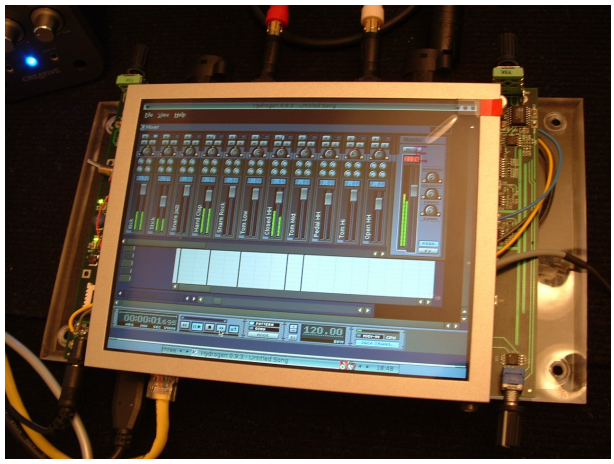

Figure 3: Berlin pre-production model, case off, running Hydrogen, 2007

A second pre-production model, featuring a more practical aluminium case, made its debut at the Podcast and New Media Expo in Ontario, California, USA, six months later.

Figure 4: Ontario model running Ardour, 2007

A short construction run of seven Ontario units had identified several ways to make the device thinner and lighter. However, the custom built hardware proved just too expensive for the product

to be sold anywhere near the target retail price of $1000 US dollars.

A meeting between Ronald Stewart and Intel resulted in the suggestion that we consider that company's Ultra Mobile PC (UMPC) platform as the basis of the Trinity product. UMPC devices are made by several OEMs, including Samsung. Samsung offers the Q1 Ultra tablet device, which was considered closest to the Trinity concept in design. It is a standard x86 device but is not much larger than a games console, being designed to be held between the user's thumbs.

The Q1 Ultra features an ultra-low-voltage Celeron CPU clocked at 900MHz, up to 2GB RAM and 40GB storage. The display is a seven inch wide-format touchscreen with 1024x600 pixels. On-board sound is provided by an Intel HDA chipset, which in theory is capable of 32-bit S/PDIF input and output. Subjectively, audio quality proved to be acceptable for headphone and loudspeaker monitoring, although phantom powered microphones were not supported. It was decided to work around this limitation by using the Mic Port Pro, a small 24-bit class-compliant USB interface with an XLR socket and phantom power. This device attaches to the base of a condensor microphone and provides a USB cable connection for the Q1 Ultra. Some condenser microphones are now being produced with USB interfaces built-in, but for now, few of these microphones support 24-bit audio sampling.

Figure 5: Transmission on the Samsung Q1 Ultra, Intel booth at CES trade show, 2008

After further software customisation, the Q1 Ultra product was launched to the public under the brand name 'Indamixx' in the summer of 2008, initially at a price of $999 US dollars for the base model. A joint marketing campaign with Samsung led to the creation of a video-based website game at http://www.mydjfantasy.com/ which includes an animated preview of the DJ application Mixxx running on the Q1 Ultra. The Indamixx gained several good reviews from the pro audio press. It also won an award for the most innovative product of the year from 'Remix' magazine (USA), against competition from some well established names in the pro audio industry.

There are likely to be significant improvements in CPU performance, RAM quantity and disc throughput of the Indamixx, as development at the ultra-mobile end of the hardware market is moving fast. A second Indamixx model, based this time on a netbook form factor with an Intel Atom CPU, was launched in December 2008. Economies of scale in netbook production, and cost savings from eliminating the touchscreen from the design, mean that the Atom version can be sold for $499 US dollars.

4 A custom distro

Daniel James met Ronald Stewart in Los Angeles during December 2006, having read the initial press release about the Trinity project on the Linuxdevices.com site. Stewart showed James the second prototype based on the CM-iGLX module and the development board, running Xubuntu. As our company had very recently released the 64 Studio 1.0 multimedia distro, we knew that we could make some major improvements to the software running on the device. It's not that there is anything wrong with Xubuntu or Ubuntu for general use, just that when creating an embedded pro audio system which departs from the desktop metaphor, some modifications are inevitably required. For one thing, the Trinity device had no QWERTY keyboard at that point, and the display is small by PC or laptop standards - the screen on the prototypes was up to eight inches diagonally, with an 800x600 resolution.

In addition, the emphasis on audio quality and the crucial requirement for glitch-free operation mean that a real-time kernel was required, which still does not come as standard with any mainstream desktop distro. At 64 Studio we'd had

lots of experience of using Ingo Molnar's RT patch sets with successive Linux kernel releases, and of using techniques like rtlimits to make sure that audio processing always gets priority. The GUI may slow down, writes to disc may be delayed, but when a musician is playing live in front of 10,000 people, the audio must not break up as system load increases. The same goes for critical live recordings, where capturing the audio a second time may be impossible.

We proposed to Stewart that we would create a custom distribution for the Trinity device. This would be based initially on the Debian 'Etch' release, using many of the same modifications that we had made for the 64 Studio distro and an existing OEM customer. In 2008, the distribution was updated against Ubuntu LTS 'Hardy' sources, to provide newer versions of core packages. The 64 Studio Platform Development Kit, implemented by lead developer Free Ekanayaka, allows multiple bespoke distros to be created from Debian and Ubuntu sources, automating many routine maintenance tasks, and sharing base packages between them.

We also suggested, given the thumb-powered interface to the Trinity hardware, that a custom GUI be designed. We chose Fluxbox as the starting point for this GUI, as we were familiar with it, it is stable, and it makes extremely light demands on system resources. We wanted to save every bogomips possible of the CPU's power for audio work; although the Geode and Celeron are surprisingly nippy for embedded CPUs consuming just a few watts of power, any single-core x86 chip running at less than 1GHz is unlikely to approach contemporary desktop performance, and so efficiencies have to be found. It's simply a waste of resources to attempt a desktop with animated cursors and rotating 3D effects on a machine of this type, since on the original Ontario baseboard there was no hardware graphics acceleration available.

The Samsung Q1 Ultra does have hardware 3D support in X.org, due to its use of an Intel video chipset. However, experiments with Compiz showed that the integration between the 3D desktop and OpenGL accelerated applications, in particular Mixxx, was poor. By the time of the Indamixx product release, we had settled on a customised version of ume-launcher from Ubuntu's Netbook Remix project as the main GUI.

The custom distribution created for the Trinity device is called Transmission. It uses the Debian installer but with all questions pre-seeded, so users don't need to know anything about Debian, Ubuntu or GNU/Linux to format and reinstall a Trinity device with a blank or corrupted hard disc. Currently, the installer defaults to the English language, a US keyboard layout and a US time zone.

There is only one user, 'trinity', which is set to log in without a password by default. The gdm package is included, and so the device can be set up for multiple users with passwords if required. Virtual keyboards aren't much use for login, since they aren't usually available until after that step. We have considered authentication through the use of encryption keys on USB memory sticks, for a more secure multi-user setup without a keyboard.

5 Application selection

At the core of the system is the jackdmp package, the forthcoming version 2 of the JACK Audio Connection Kit. The jackdmp server runs by default after the Indamixx is switched on by the user, with a small applet in the GUI indicating JACK status. Alongside, another applet for Network Manager shows Internet connectivity status.

Ardour provides the recording, editing and mixing features that the Trinity needs, but it was originally designed for large, multi-monitor displays of at least 1280x1024 pixels apiece. Audacity has a much simpler and smaller GUI, but its JACK support is provided via the PortAudio library, and is therefore not as 'native' to JACK, compared with Ardour. Ardour also has the benefit of real-time DSP effects in its mixer, provided by Free Software plugins in the LADSPA format. Having established that Ardour would run on a machine with such modest disc throughput, due to the use of a relatively slow miniature drive, Ekanayaka modified the GTK+ GUI so that only the most frequently used controls were visible. This meant that Ardourino, as we code-named the new version, could now fit into a small display.

The amSynth application, which models an analogue two-oscillator synthesiser in software, was initially included in the distro, but had to be removed after the switch to jackdmp, which amSynth does not support. ZynAddSubFX and ALSA Modular Synth applications provide alternative synths. The Seq24 sequencer was added to the distro to provide a lightweight MIDI sequencing tool.

The Hydrogen drum sequencer was also included, with a selection of sampled and synthesised drum kits provided by Indamixx beta testers in the Los Angeles area. Hydrogen, AMS and ZynAddSubFX can respond to MIDI sent by external devices, such as the portable, battery-operated USB MIDI keyboards from Novation and other manufacturers.

Mixxx, the DJ application, works particularly well with the touchscreen of the Q1 Ultra. A Mixxx skin had to be adapted to fit the widescreen aspect ratio of the display. Mixxx also supports hardware control surfaces produced by Hercules and Behringer. The proprietary EnergyXT application has been added to Indamixx to address the need for an all-in-one sequencing and recording application.

One application that takes particular advantage of the wireless and Ethernet connectivity of the Trinity is idjc, the Internet DJ Console. At first glance this application is a two-playlist media player with a crossfader, in the manner of other DJ software, but it is an Icecast/Shoutcast streaming client too. It contains many other features useful to radio-style broadcasting, such as jingle buttons, microphone compression, and even VoIP integration so that you can host your own phone-in show. The idjc application means that you can hook up an Indamixx with a microphone, and broadcast live from anywhere with a Wi-Fi access point available - without being rooted to the spot by cables.

There are several other applications included in Transmission with mobile media production in mind. The Gnome CD Master application can burn audio files to CD, if a USB CD writer is connected. It's one of the less well known Free Software applications for preparing CDs, but it does have unique features for audio mastering - like being able to place track markers directly on the waveform by ear, which is very useful for dividing up continuous live recordings. On the networking side, gFTP and the Firefox browser are used for file uploads. Ubuntu LTS package compatibility means that many other programs can be added to the Transmission distro, using any Ubuntu mirror.

A retro-fit version of Transmission, which will allow netbook owners to re-purpose their devices into mobile DAWs, is currently in preparation. The Transmission distro project is organised much like any other Free Software development effort, with a Trac-based website and a mailing list. Community involvement in the project is welcomed, and because the distro is x86 compatible, you don't need any special hardware to take part. Any 32-bit or 64-bit x86 computer will serve as a development platform. Since the distro is Debian/Ubuntu based, you can update the machine easily with APT.

We look forward to hearing your ideas for the Indamixx, so we can make this product not only the first mobile audio computer powered by GNU/Linux and Free Software, but also the 'must-have' platform for musicians, sound engineers and broadcasters all over the world.

References

Trinity homepage
http://www.trinityaudiogroup.com/

Indamixx website
http://www.indamixx.com/

Transmission development site
http://trac.64studio.com/trinity/

Ardour
http://www.ardour.org/

Hydrogen drum machine
http://www.hydrogen-music.org/

Internet DJ Console
http://www.onlymeok.nildram.co.uk/